

# Global Availability of NVIDIA AI Enterprise Makes AI Accessible for Every Industry

## NVIDIA Software Gives Hundreds of Thousands of Enterprises Worldwide Ability to Run AI on VMware vSphere and Industry-Standard Servers

NVIDIA today announced the availability of [NVIDIA AI Enterprise](#), a comprehensive software suite of AI tools and frameworks that enables the hundreds of thousands of companies running VMware vSphere to virtualize AI workloads on [NVIDIA-Certified Systems](#)<sup>™</sup>.

Leading manufacturers Atos, Dell Technologies, GIGABYTE, Hewlett Packard Enterprise, Inspur, Lenovo and Supermicro are offering NVIDIA-Certified Systems optimized for AI workloads on VMware vSphere with NVIDIA AI Enterprise. Separately, Dell Technologies today announced Dell EMC VxRail as the first hyperconverged platform to be qualified as an NVIDIA-Certified System for NVIDIA AI Enterprise.

To help teams of data scientists run their AI workloads most efficiently, [Domino Data Lab](#) today announced it is validating its Domino Enterprise MLOps Platform with NVIDIA AI Enterprise, which runs on mainstream NVIDIA-Certified Systems.

“The first wave of AI has been powered by specialized infrastructure that focused adoption on industry pioneers,” said Manuvir Das, head of Enterprise Computing at NVIDIA. “Today is the beginning of a new chapter in the age of AI, as NVIDIA software brings its transformative power within reach for enterprises around the world that run their workloads on VMware with mainstream data center servers.”

“As AI applications become critical, customers want to run them on their enterprise infrastructure for manageability, scalability, security and governance,” said Krish Prasad, senior vice president and general manager of the Cloud Platform Business Unit at VMware. “Running NVIDIA AI Enterprise on VMware vSphere delivers a certified, end-to-end AI-Ready enterprise platform that’s easy to deploy and operate.”

“Partnering closely with NVIDIA, we’re deepening our product integrations by enabling the Domino Enterprise MLOps Platform to run with a broader range of NVIDIA GPUs and validating it for NVIDIA AI Enterprise,” said Nick Elprin, CEO and co-founder of Domino Data Lab. “This new offering will help hundreds of thousands of enterprises accelerate data science at scale.”

### Customers Simplify, Scale with NVIDIA AI Enterprise

Dozens of automotive, education, finance, healthcare, manufacturing and technology companies worldwide are among the early adopters using NVIDIA AI Enterprise. Many are midsize companies that can now develop a broad range of applications using the world’s most widely used servers to deploy and scale data science, conversational AI, computer vision, recommender systems and more.

Among the earliest to use NVIDIA AI Enterprise is Cerence Inc., a leading provider of conversational AI for automotive and mobility markets with nearly 400 million Cerence-powered vehicles shipped worldwide. The company is using AI Enterprise to develop intelligent in-car assistants and digital co-pilots.

Additionally, University of Pisa, an Italian public research university, is supporting HPC and AI training across multiple disciplines to advance scientific studies with the NVIDIA software.

“NVIDIA AI Enterprise allowed us to expand our support for our researchers and students who utilize data analytics and AI deep learning and machine learning, while making these applications easier to deploy and manage,” said Maurizio Davini, chief technology officer at the University of Pisa. “Our testing has shown that these latest collaborations between NVIDIA and VMware deliver the full potential of our GPU-accelerated virtualized infrastructure at near bare-metal speeds.”

### Powerful AI Performance on Mainstream Servers

NVIDIA AI Enterprise enables IT professionals that use VMware vSphere to run traditional enterprise applications to easily and cost-effectively support AI workloads while using the same tools they use to manage large-scale data centers and hybrid clouds.

NVIDIA-Certified Systems from Atos, [Dell Technologies](#), [GIGABYTE](#), [Hewlett Packard Enterprise](#), [Inspur](#), [Lenovo](#) and Supermicro for NVIDIA AI Enterprise feature a wide range of NVIDIA GPUs, including the [A100](#), [A30](#), [A40](#), [A10](#) and [T4](#). These mainstream accelerated systems provide customers with a broad selection of options for scale-out, multinode, AI application performance on vSphere that is virtually indistinguishable from bare-metal servers.

### Availability

NVIDIA AI Enterprise is now available from NVIDIA channel partners worldwide, including Atea, [Carahsoft](#), [Computacenter](#), [Insight Enterprises](#), [NTT](#), [SoftServe](#) and SVA System Vertrieb Alexander GmbH.

Subscription licenses start at \$2,000 per CPU socket for one year and include Business Standard Support (five days a week, nine hours a day). Perpetual licenses are \$3,595 and require additional support purchase. Customers can upgrade to Business Critical Support for 24x7 access to NVIDIA AI expertise.

#### **About NVIDIA**

[NVIDIA](#)'s (NASDAQ: NVDA) invention of the GPU in 1999 sparked the growth of the PC gaming market and has redefined modern computer graphics, high performance computing and artificial intelligence. The company's pioneering work in accelerated computing and AI is reshaping trillion-dollar industries, such as transportation, healthcare and manufacturing, and fueling the growth of many others. More information at <https://nvidianews.nvidia.com/>.

Certain statements in this press release including, but not limited to, statements as to: the benefits, impact, performance, features, pricing, and availability of our products and services, including NVIDIA AI Enterprise; leading manufacturers offering NVIDIA-Certified Systems optimized for AI workloads on VMware vSphere with NVIDIA AI Enterprise; the impact of our partnership with Domino Data Lab; NVIDIA software bringing its transformative power within reach for enterprises around the world; customers adopting and using NVIDIA AI Enterprise; NVIDIA-Certified Systems providing customers with a broad selection of options for scale-out, multinode, AI application performance; and the availability of NVIDIA AI Enterprise from our channel partners worldwide are forward-looking statements that are subject to risks and uncertainties that could cause results to be materially different than expectations. Important factors that could cause actual results to differ materially include: global economic conditions; our reliance on third parties to manufacture, assemble, package and test our products; the impact of technological development and competition; development of new products and technologies or enhancements to our existing product and technologies; market acceptance of our products or our partners' products; design, manufacturing or software defects; changes in consumer preferences or demands; changes in industry standards and interfaces; unexpected loss of performance of our products or technologies when integrated into systems; as well as other factors detailed from time to time in the most recent reports NVIDIA files with the Securities and Exchange Commission, or SEC, including, but not limited to, its annual report on Form 10-K and quarterly reports on Form 10-Q. Copies of reports filed with the SEC are posted on the company's website and are available from NVIDIA without charge. These forward-looking statements are not guarantees of future performance and speak only as of the date hereof, and, except as required by law, NVIDIA disclaims any obligation to update these forward-looking statements to reflect future events or circumstances

© 2021 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, and NVIDIA-Certified Systems are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. VMware and vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and other jurisdictions. All other trademarks and copyrights are the property of their respective owners. Features, pricing, availability, and specifications are subject to change without notice.

Shannon McPhee  
+1-310-920-9642  
[smcphee@nvidia.com](mailto:smcphee@nvidia.com)